MULTIDIMENSIONAL DATABASE STORAGE AND RETRIEVAL SYSTEM

ABSTRACT OF THE DISCLOSURE

In a multidimensional database, an aggregation operation is performed in an optimal manner by storing the values included in the aggregation operation on the same disk page. A sparsity manager determines aggregate values that are computed from other data values during the aggregation operation. Each aggregate value is associated with one or more data values that are used during the aggregation operation to compute the aggregate value. The sparsity manager stores the associated data values in proximity to each other, such as on the same disk page, so that multiple disk page fetches may not be required for the same set of data values during the aggregation operation. The data values used in the aggregation operation can therefore be fetched once from a disk page, and thereafter are found in memory, such as on a cache page corresponding to the disk page. In this manner, multiple fetches for the same disk page during the aggregation operation are avoided.